

EXPLORING HOOVER AND PEREZ'S EXPERIMENTAL DESIGNS USING GLOBAL SENSITIVITY ANALYSIS

WILLIAM BECKER[°], PAOLO PARUOLO[°], AND ANDREA SALTELLI[°]

ABSTRACT. This paper investigates variable-selection procedures in regression that make use of global sensitivity analysis. The approach is combined with existing algorithms and it is applied to the time series regression designs proposed by Hoover and Perez. A comparison of an algorithm employing global sensitivity analysis and the (optimized) algorithm of Hoover and Perez shows that the former significantly improves the recovery rates of original specifications.

1. INTRODUCTION

Model selection in regression analysis is a central issue, both in theory and in practice. A partial list of statistical fields with a non-empty intersection with model selection includes multiple testing, see e.g. [Romano and Wolf \(2005\)](#) and [Bittman et al. \(2009\)](#), pre-testing, see [Leeb and Poetscher \(2006\)](#), information criteria, see [Hjort and Claeskens \(2003\)](#) and [Liu and Yang \(2011\)](#), model selection based on Lasso, see e.g. [Brunea \(2008\)](#), model averaging, see [Claeskens and Hjort \(2003\)](#), stepwise regression, see [Miller \(2002\)](#), risk inflation in prediction, see [Foster and George \(1994\)](#), directed acyclic graphs and causality discovery, see e.g. [Freedman and Humphreys \(1999\)](#).¹

Model choice is also of primary concern in many areas of applied econometrics, as witnessed for example by the literature on growth regression, see e.g. [Sala-i-Martin \(1997\)](#). Controlling for the right set of covariates is central in the analysis of policy impact evaluations; this is embodied in the assumption of unconfoundedness, see e.g. [Imbens and Wooldridge \(2009\)](#). In economic forecasting, model selection is the main alternative to model averaging, see e.g. [Hjort and Claeskens \(2003\)](#).

The analysis of the effects of pre-testing on parameter estimation has a long tradition in econometrics, see [Danilov and Magnus \(2004\)](#) for a recent account; in this context [Magnus and](#)

Date: January 23, 2014.

Key words and phrases. Model selection, Sensitivity Analysis, Simulation, Monte Carlo.

[°] Econometrics and Applied Statistics Unit, European Commission, DG JRC IPSC, Via E.Fermi 2749, I-21027 Ispra (VA), Italy; emails: william.becker@jrc.ec.europa.eu, paolo.paruolo@jrc.ec.europa.eu, andrea.saltelli@jrc.ec.europa.eu.

¹Model selection is also associated with current rules of thumb on the maximum number of regression parameters to consider. This literature appears to have been initiated by [Freedman \(1983\)](#), who considered the case of a first screening regression with 50 regressors and 100 data points, where regressors that are significant at 25% significance level are kept in a second regression. Freedman showed that the second regression is troublesome when one acts as if the screening regression had not been performed and the ratio of number of observations to number of regressors in the screening regression is kept in a fixed proportion as the number of observations diverges. This study was followed by [Freedman and Pee \(1989\)](#), [Freedman et al. \(1992\)](#), who defined the rule of thumb that the ratio of the number of observations per regressor should be at least equal to 4; this rule is included in [Harrell \(2001\)](#), who suggested to have it at least equal to 10.

Durbin (1999) and co-authors proposed the weighted average least squares estimator (WALS), and compared it with model averaging for growth empirics, see Magnus et al. (2010).

Model selection is a major area of investigation also in time-series econometrics, see e.g. Phillips (1997, 2003). The so-called London School of Economics (LSE) methodology has played a prominent role in this area, advocating the *general-to-specific* (GETS) approach to model selection, see Hendry and Krolzig (2005), Castle et al. (2011) and references therein. In a widely cited paper, Hoover and Perez (1999) (hereafter HP) ‘mechanized’ – i.e. translated – the GETS approach into an algorithm for model selection; they then tested the performance of the HP algorithm on a set of time-series regression designs, constructed along the lines of Lovell (1983).

Model selection is also related to the issue of regression coefficients’ robustness (i.e. lack of sensitivity) to the omission/inclusion of additional variables. Leamer (1983) has proposed extreme bound analysis, i.e. to report the range of possible parameter estimates of the coefficient of interest when varying the additional regressors included in the analysis, as an application of sensitivity analysis to econometrics. Other applications of sensitivity analysis to econometrics include the local sensitivity to model misspecification developed in Magnus and Vasnev (2007) and Magnus (2007).²

Sensitivity analysis originated in the natural sciences, and it is generally defined as ‘the study of how the uncertainty in the output of a mathematical model or system (numerical or otherwise) can be apportioned to different sources of uncertainty in its inputs’, see Saltelli (2002). Deterministic models based on knowledge of the physical laws governing the system are usually fruitfully applied in physical sciences. Box and Draper (2007) advocate their use (in combination with statistical models) because they (i) contribute to the scientific understanding of the phenomenon under study, (ii) provide a better basis for extrapolation with respect to empirical models (iii) they tend to be parsimonious (i.e. frugal) in the use of parameters. The combined use of deterministic and stochastic models is also advocated in other non-experimental fields, such as in environmental modeling, see Young et al. (1996).

Despite several uses of sensitivity in econometrics, the present authors are not aware of systematic applications of the techniques of Global Sensitivity Analysis, GSA, see Saltelli et al. (1993), to the problem of model selection in regression. The present paper attempts a first experimental exploration of the possible application of GSA to model selection in time-series regression. Here we aim to answer the question: “Can GSA methods help in model selection in practice?”. This question is answered in the affirmative, using the ‘total sensitivity index’ to rank regressors’ importance in order to construct relevant subsets of models.

For simplicity and in order to increase replicability of our exploration, we have chosen to compare new tools and old via simple Monte Carlo (MC) methods.³ We have chosen to replicate the original search algorithm in HP as a benchmark, and to compare our GSA algorithm with

²They show that local sensitivity measures provide complementary information with respect to standard diagnostic tests for misspecification, i.e. that the two types of statistics are asymptotically independent. In SA a local measure of sensitivity is one focused on a precise point in the space of the input factor, a.g. a partial derivative of the output versus the input. With a global measure of sensitivity the influence of a given input on the output is averaged both on the distribution of the input factor itself and on the distributions of all the remaining factors, see Saltelli et al. (1993).

³See however analytical results on the properties of the GSA-based algorithm in Appendix A.

the original HP algorithm. Because the purpose of the paper is to investigate the contribution of GSA methods, we have abstained from the implementation of any other algorithms except the original – albeit optimized – HP algorithm as a benchmark.

The choice of the design of experiments in HP reflects the current practice in single-equation, time-series econometric models; these consist in a possibly dynamic regression equation with exogenous variables, where the exogenous variables are fixed across experiments and are taken from real-world, stationary, macroeconomic time series. While HP's designs are supposed to represent prototypical configurations in time-series econometrics, they contain by construction only a small subset of possible situations encountered in econometric time-series applications. As such, it is like a single planet in a galaxy.

As forbidding as the exploration of a galaxy is (at least with current means), so is the attempt to investigate all regression designs. In this paper we have hence decided to explore a limited part of this galaxy – a single planet – namely HP's experimental designs.

Several papers appear to have applied other methods to HP's designs, see [Hendry and Krolzig \(1999\)](#), [Castle et al. \(2011\)](#). The choice of HP's designs and of the HP algorithm as benchmark allows to compare performances in our paper with others reported in the literature. The designs in HP's designs include data generating processes (DGPs) of varying degree of difficulty (for model search algorithms) and a single sample size of 139 time periods, close to the ones available in typical macroeconomic applications with quarterly data.

The features of HP's designs prompt a number of considerations. First, because sample size is limited and fixed, consistency of model-selection algorithms cannot be the sole performance criterion. In this light, it would be helpful to be able to describe the complete finite sample properties of model-selection algorithms for HP's designs; the MC approach taken in the paper allows to do this.

Secondly, some of the DGPs in HP's designs are characterized by a low signal-to-noise ratio for some coefficients; we call the corresponding regressors 'weak'. This situation makes it very difficult for statistical procedures to discover if the corresponding regressors should be included or not. This raises the question of how to measure selection performance in this context.

In this paper we observe that, in the case of weak regressors, one can measure performance of model-selection algorithms also with respect to a simplified DGP, which contains the subset of regressors with sufficiently high signal-to-noise ratio; we call this the 'Effective DGP', EDGP. In this paper we make the definition of the EDGP operational using the 'parametricness index' recently introduced in [Liu and Yang \(2011\)](#).

Overall, results point to the possible usefulness of GSA methods in model selection algorithms. In particular, GSA methods seem to complement existing approaches, as they give a way to construct viable search paths (via ordering of regressors) that are complementary to the ones based on t ratios. When comparing optimized algorithms, the GSA method appears to be able to reduce the failure rate in recovering the EDGP from 5% to 1% approximately – a five-fold reduction. When some of the regressors are weak, the recovery of exact DGPs does not appear to be improved by the use of GSA methods.

Selection of a subset of all possible models from the space of all submodels is one of the critical aspects also for model selection based on information criteria, see Section 5.2. in [Hansen \(1999\)](#).

A similar remark applies for multi-model inference procedures, see e.g. [Burnham and Anderson \(2002\)](#). The results obtained in this paper show that GSA methods have potential to make these methods operational. Due to space limitations, we leave the analysis of these extensions to future research.

The rest of the paper is organized as follows. In [Section 2](#) we define the problem of interest and introduce HP's data generating processes and the HP algorithm. [Section 3](#) defines the tools from GSA used in the paper, while [Section 4](#) presents the GSA algorithm. Results are reported in [Section 5](#), and [Section 6](#) concludes. Large-sample properties of the orderings based on the GSA algorithm are discussed in [Appendix A](#). A discussion about the identifiability of DGPs and the definition of EGDP is reported in [Appendix B](#). Finally, this paper follows the notational conventions in [Abadir and Magnus \(2002\)](#).

2. PROBLEM DEFINITION

This section presents the setup of the problem and describes the design of experiments in HP, as well as their algorithm.

2.1. Model selection in regression. Let n be the number of data points and p the number of regressors in a standard multiple regression model of the form

$$\mathbf{y} = \mathbf{X}_1\beta_1 + \dots \mathbf{X}_p\beta_p + \varepsilon = \mathbf{X}\beta + \varepsilon \quad (1)$$

where $\mathbf{y} = (y_1, \dots, y_n)'$ is $n \times 1$, $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_p)$ is $n \times p$, $\mathbf{X}_i := (x_{i,1}, \dots, x_{i,n})'$ is $n \times 1$, $\beta = (\beta_1, \dots, \beta_p)'$ is $p \times 1$ and ε is a $n \times 1$ Gaussian random vector with distribution $N(\mathbf{0}, \sigma^2 \mathbf{I}_n)$. The symbol $'$ indicates transposition.

Let Γ be the set of all $p \times 1$ vectors of indicators $\gamma = (\gamma_1, \dots, \gamma_p)'$, with $\gamma_i = 0$ or 1 for $i = 1, \dots, p$, i.e. $\Gamma = \{0, 1\}^p$. A submodel of [\(1\)](#) (or one specification) corresponds to one vector $\gamma \in \Gamma$, where $\gamma_i = 0$ (respectively 1) indicates that β_i is to be estimated equal to 0 (respectively unrestrictedly). Note that there are 2^p different specifications, i.e. γ vectors in Γ . When $p = 40$ as in HP's designs, the number of specifications $2^p \approx 1.0995 \cdot 10^{12}$ is very large.

In the following we indicate by $\beta_0 = (\beta_{0,1}, \dots, \beta_{0,p})'$ the true value of β . Define also $\gamma_0 = (\gamma_{0,1}, \dots, \gamma_{0,p})'$ with $\gamma_{0i} = 1(\beta_{0,i} \neq 0)$, where $1(\cdot)$ denotes the indicator function and $\beta_{0,i}$ are the true parameters in [\(1\)](#). The vector of indicators γ_0 defines the smallest true submodel; this is called the *Data Generating Process* (DGP) in the following.

The least squares estimator of β in model γ can be written as follows:

$$\hat{\beta}_\gamma = (\mathbf{D}_\gamma \mathbf{X}' \mathbf{X} \mathbf{D}_\gamma)^+ \mathbf{D}_\gamma \mathbf{X}' \mathbf{y}, \quad (2)$$

where $\mathbf{D}_\gamma = \text{diag}(\gamma)$ is the $p \times p$ matrix with diagonal elements γ and \mathbf{A}^+ indicates the Moore-Penrose generalized inverse of the matrix \mathbf{A} . The non-zero elements in $\hat{\beta}_\gamma$ correspond to the least squares estimates in the submodel which includes only regressors \mathbf{X}_i for which $\gamma_i = 1$.

The case of γ equal to $\mathbf{1}$, a vector with all 1s, is called the *General Unrestricted Model*, the GUM in HP. The problem of interest is, given the observed data, to find γ_0 , i.e. to identify the DGP.⁴

⁴All empirical models are assumed to contain the constant; this is imposed implicitly by de-meaning the \mathbf{y} and \mathbf{X}_i vectors. Hence in the following, the 'empty set of regressors' refers to the regression model with only the constant.

In this paper we assume that the model is correctly specified, i.e. that γ_0 is an element of Γ . This is a common hypothesis in the regression literature. In econometrics this assumption appears to be more questionable, because of the possibility of relevant omitted variables. However, we maintain it here for reasons of simplicity.

2.2. HP's designs. HP's designs are constructed as follows. Following [Lovell \(1983\)](#), HP chose a set of 18 major US quarterly macroeconomic variables. Only two variables considered in [Lovell \(1983\)](#) were discarded in HP, namely the linear trend and the 'potential level of GNP in \$1958', because they were no longer relevant or available. Unlike in [Lovell \(1983\)](#), HP applied 0, 1 or 2 differences to the data; the order of differencing was selected by HP in order to obtain stationary variables according to standard unit root tests, see their Table 1.

The values of these (differenced) 18 major US quarterly macroeconomic series are then fixed in HP's designs; they are here indicated as x_{it}^* , where $t = 1, \dots, n$ indicates quarters and $i = 1, \dots, k$, with $k = 18$ indexes variables. The values of y_t were then generated by the following scheme

$$y_t = \sum_{i=1}^k \beta_i^* x_{it}^* + u_t \quad u_t = \rho u_{t-1} + \varepsilon_t, \quad (3)$$

where ε_t are i.i.d. $N(0, \sigma^2)$. Here β_i^* for $i = 1, \dots, k$ and σ^2 are known constants, which define the DGP. In practice ε_{ts} are simulated using a computer random number generator, u_t is then calculated as an autoregressive series of order 1, AR(1), with coefficient ρ . u_t is then fed into the equation for y_t , where x_{it}^* are kept fixed and do not change across replications.

It is useful to express (3) as a special case of (1). To this end one can substitute $(y_t - \sum_{i=1}^k \beta_i^* x_{it}^*)$ in place of u_t in the dynamic equation of u_t ; one hence finds the following equivalent representation of the DGP

$$y_t = \rho y_{t-1} + \sum_{i=1}^{2k} \beta_i x_{it} + \varepsilon_t \quad (4)$$

where $\beta_i = \beta_i^*$ and $x_{it} = x_{it}^*$ for $i = 1, \dots, k$ while $\beta_i = -\rho \beta_i^*$ and $x_{it} = x_{it-1}^*$ for $i = k+1, \dots, 2k$. This representation is in the form (1), and the parameters can be estimated as in (2).

Regressions in HP were performed setting the elements $x_{i,t}$ in column \mathbf{X}_i equal to variable x_{it} from (4), for $i = 1, \dots, 2k$ with $2k = 36$, and setting the elements $x_{i,t}$ of the remaining columns \mathbf{X}_i for $i = 2k+1, \dots, p$, i.e. from 37 to 40, equal to the first, second, third and fourth lag of y_t . Therefore, 4 lags were always considered in estimation (even if only one lag was possibly present under the DGP), and the only part of the \mathbf{X} that changes across replications is the last 4 columns.

HP defined 11 experimental designs (DGPs) by choosing values for the parameters ρ , β_i^* and σ_ε^2 . Table 1 summarizes the chosen parameter values. The choice of these values was made to reflect the coefficient estimates obtained on US data, using personal consumption expenditure as dependent variable, following the rationale in [Lovell \(1983\)](#). Because they were chosen as explanatory variables for a consumption equation, not all the macroeconomic time series were included in the DGP; in particular only (the second differences of the) Government purchases on goods and services G and the (first differences of the) $M1$ monetary aggregate, and their respective first lags, were included in the designs.

DGP	1	2	3 [‡]	4	5	6	6A	6B	7	8	9
coefficients in DGP											
y_{t-1}		0.75	0.395						0.75	0.75	0.75
y_{t-2}			0.3995								
G_t					-0.046	-0.023	-0.32	-0.65		-0.046	-0.023
G_{t-1}										0.00345	0.01725
$M1_t$				1.33		0.67	0.67	0.67	1.33		0.67
$M1_{t-1}$									-0.9975		-0.5025
σ_ε	130	85.99	0.00172	9.73	0.11	4.92	4.92	4.92	6.73	0.073	3.25

TABLE 1. DGPs design. y_{t-j} indicates lags of the dependent variable, G_{t-j} denotes (lags of) second differences of government purchases of goods and services and $M1_{t-j}$ indicates (lags of) first differences of M1.

[‡]: in DGP 3 the regression analysis is performed on $y_t^* = \exp(y_t)$, where y_t is simulated as in (4).

2.3. HP algorithm. HP proposed an algorithm that aims to provide a close approximation to a subset of what practitioners of the LSE approach actually do. Here we follow Hansen (1999) in his description of the HP algorithm.

The HP algorithm can be described by a choice of a triplet (R, f, Γ_s) composed of (i) a test procedure R , (ii) a measure of fit f and (iii) a subset Γ_s of all models Γ , $\Gamma_s \subseteq \Gamma$. For any model γ , the test procedure R is defined as

$$R(\gamma) = 1(\min_{1 \leq \ell \leq v} p_\ell \leq \alpha) \quad (5)$$

where p_ℓ are the p -values of v specification tests and α is the chosen significance level. Note that $R(\gamma) = 0$ when all v tests do not reject the null, which corresponds to the hypothesis of correct specification and/or constant parameters.⁵

HP's measure of fit f is based on the least-square estimate of σ^2 , the regression variance, which equals $\tilde{\sigma}_\gamma^2 := \frac{1}{n-k_\gamma} \hat{\varepsilon}_\gamma' \hat{\varepsilon}_\gamma$, where k_γ and $\hat{\varepsilon}_\gamma$ are the number of regressors and the residuals in model γ . HP's measure of fit is $f(\gamma) = \tilde{\sigma}_\gamma$, which should be minimized. Finally the subset Γ_s is selected recursively, going from general to specific models, starting from the GUM, $\gamma = \mathbf{z}_p$; the recursion continues as long as $R(\gamma) = 0$. Details on HP's choice of Γ_s are given in the next section.

Overall the HP algorithm selects a model $\hat{\gamma}$ as the preferred model using the rule

$$\hat{\gamma} = \arg \min_{\gamma \in \Gamma_s: R(\gamma)=0} f(\gamma).$$

The above description shows that the HP algorithm depends on α , which is a tuning parameter, as well as on the choice of specific path Γ_s . For large n , Hansen (1999) noted that $\hat{\gamma}$ corresponds

⁵The tests are the following: (1) Jarque Bera test for normality of residuals; (2) Breusch Pagan residual autocorrelation tests; (3) Engle's ARCH test on residuals; (4) Chow sample-split parameter stability tests; (5) Chow out-of-sample stability test using the first 90% of observations versus the last 10%; (6) F test of the restrictions imposed by model γ versus the GUM. The tests are performed on the first 90% of observations during the search.

approximately to minimizing the information criterion⁶ $HP(\gamma) = \log \hat{\sigma}_\gamma^2 + k_\gamma/n$, where $\hat{\sigma}_\gamma^2 := \frac{1}{n} \hat{\varepsilon}_\gamma' \hat{\varepsilon}_\gamma$ is the ML estimator of σ^2 . This differs from Akaike's Information Criterion $AIC(\gamma) = \log \hat{\sigma}_\gamma^2 + 2k_\gamma/n$ and from the Bayesian Information criterion of Schwarz $BIC(\gamma) = \log \hat{\sigma}_\gamma^2 + k_\gamma \log(n)/n$ by the different choice of penalty term.⁷

2.4. A subset of models. The number of models in Γ is too large to visit all submodels; hence any selection method needs to select at most a subset Γ_s of Γ . This is a critical aspect of the HP algorithm, as well as of any selection method based e.g. on information criteria, see Section 5.2. in Hansen (1999) and Burnham and Anderson (2002).

In particular, HP select a subset Γ_s as follows. All paths start from the GUM regression, and the regressors are ranked in ascending order according to their t -statistics. The 10 lowest variables in this list are then candidates for elimination; this starts an iterative elimination path. Each candidate model γ_* then becomes the current specification provided $R(\gamma_*) = 0$. In this stage, the first 90% of the observations are used in the specification tests. Each search is terminated when for any choice of regressor the test R rejects.

At this final stage, the HP algorithm reconsiders all the observations in a 'block search'; this consists in considering the joint elimination of all the regressors with an insignificant t -statistics. If the R tests for the block search does not reject, the resulting model becomes the terminal specification. Otherwise, the specification that entered the final stage becomes the terminal specification. Once all 10 search paths have ended in a terminal specification, the final specification is the one among these with lowest $f(\gamma) = \tilde{\sigma}_\gamma$.

This paper gives a contribution on the selection of Γ_s , by defining a GSA-based ordering of regressors.

2.5. Measures of performance. The performance of algorithms was measured by HP via the number of times the algorithm selected the DGP as a final specification. Here we describe measures of performance similar to the ones in HP, as well as additional ones proposed in Castle et al. (2011).

Recall that γ_0 is the true set of included regressors and let $\hat{\gamma}_j$ indicate the one produced by a generic algorithm in replication $j = 1, \dots, N_R$. Here we define r_j to be number of correct inclusions of components in vector $\hat{\gamma}_j$, i.e. the number of regression indices i for which $\hat{\gamma}_{j,i} = \gamma_{0,i} = 1$, $r_j = \sum_{i=1}^p 1(\hat{\gamma}_{j,i} = \gamma_{0,i} = 1)$. Similarly, we let r_0 indicate the number of true regressors.

We then define the following exhaustive and mutually exclusive categories of results:

C_1 : exact matches;

C_2 : the selected model is correctly specified, but it is larger than necessary, i.e. it contains all relevant regressors as well as irrelevant ones;

C_3 : the selected model is incorrectly specified (misspecified), i.e. it lacks relevant regressors.

C_1 matches correspond to the case when $\hat{\gamma}_j$ coincides with γ_0 ; the corresponding frequency C_1 is computed as $C_1 = \frac{1}{N_R} \sum_{j=1}^{N_R} 1(\hat{\gamma}_j = \gamma_0)$. The frequency of C_2 cases is given by $C_2 =$

⁶Here the only approximation involved in the large T argument is $\log(1 + k_\gamma/T) \approx k_\gamma/T$.

⁷Remark that information criteria are equivalent to LR testing with a tunable significance level; see for instance Poetscher (1991).

$\frac{1}{N_R} \sum_{j=1}^{N_R} 1(\hat{\gamma}_j \neq \gamma_0, r_j = r_0)$. Finally, C_3 cases are the residual category, and the corresponding frequency is $C_3 = 1 - C_1 - C_2$.⁸

The performance can be further evaluated through measures taken from [Castle et al. \(2011\)](#), known as potency and gauge. First the retention rate \tilde{p}_i of the i -th variable is defined as, $\tilde{p}_i = \frac{1}{N_R} \sum_{j=1}^{N_R} 1(\hat{\gamma}_{j,i} = 1)$. Then, potency and gauge are defined as follows:

$$\text{potency} = \frac{1}{r_0} \sum_{i: \beta_{0,i} \neq 0} \tilde{p}_i, \quad \text{gauge} = \frac{1}{p - r_0} \sum_{i: \beta_{0,i} = 0} \tilde{p}_i,$$

where r_0 indicates the number of true regressors in the DGP.

Potency therefore measures the average frequency of inclusion of regressors belonging to the DGP, while gauge measures the average frequency of inclusion of regressors not belonging to the DGP. An ideal performance is thus represented by a potency value of 1 and a gauge of 0.

In calculating these measures, HP chose to discard MC samples for which a preliminary application of the battery of tests defined in (5) reported a rejection.⁹ We call this choice ‘pre-search elimination’ of MC samples.

2.6. Benchmark. In this paper we take the performance of HP’s algorithm as a benchmark. The original `MATLAB` code for HP designs and the HP algorithm was downloaded from HP’s home page.¹⁰ The original scripts were then updated to run on the current version of `MATLAB`. A replication of the results in Tables 4, 6 and 7 in HP is reported in the first panel of Table 2, using a nominal significance level of $\alpha = 1\%, 5\%, 10\%$ and $N_R = 10^3$ replications. The results do not appear to be significantly different from the ones reported in HP.

We then noted an incorrect coding in the original HP script for the generation of the AR series u_t in eq. (3), which produced simulations of a moving average process of order 1, MA(1), with MA parameter 0.75 instead of an AR(1) with AR parameter 0.75.¹¹ We hence modified the script to produce u_t as an AR(1) with AR parameter 0.75; we call this the ‘modified script’. Re-running the experiments using this modified script we obtained the results in the second panel in Table 2; for this set of simulations we used $N_R = 10^4$ replications. Comparing the first and second panels in the table for the same nominal significance level α , one observes a significant increase in C_1 catches in DGP 2 and 7. One reason for this can be that when the

⁸ C_1 corresponds to Category 1 in HP; C_2 corresponds to Category 2+Category 3–Category 1 in HP; finally C_3 corresponds to Category 4 in HP.

⁹We found the empirical percentage of samples that were discarded in this way was proportional to the significance level α . This fact, however, did not influence significantly the number of C_1 catches. We hence decided to let the HP procedure discard sample as in the original version. For the GSA algorithm we did not apply any pre-search elimination.

¹⁰<http://www.csus.edu/indiv/p/perezs/Data/data.htm>

¹¹This means that the results reported in HP for DGP 2, 3, 7, 8, 9 concern a design in which the entertained model is misspecified. The MA process can be inverted to obtain a AR(∞) representation; substituting from the y_t equation as before, one finds that the DGP contains an infinite number of lags on the dependent variable and of the x_{it}^* variables, with exponentially decreasing coefficients. The entertained regression model with 4 lags on the dependent variable and 2 lags on the x_{it}^* variables can be considered an approximation to the DGP.

	original script			modified script		
DGP	$\alpha = 0.01$	0.05	0.1	0.01	0.05	0.1
1	81.1	28.6	6.8	79.3	30.0	7.3
2	1.2	0.0	0.0	77.0	27.0	6.9
3	71.4	27.2	9.1	71.7	27.3	6.9
4	78.2	31.2	6.4	81.8	31.1	7.0
5	80.9	30.1	7.4	80.7	29.9	6.4
6	0.2	1.0	0.7	0.4	0.5	0.6
6A	68.0	27.8	7.8	70.6	27.8	7.6
6B	80.8	30.7	7.8	81.1	31.4	8.0
7	23.6	4.7	0.3	75.7	26.7	7.6
8	80.6	31.0	8.0	79.2	30.3	9.4
9	0.1	0	0	0	0	0

TABLE 2. Percentages of Category 1 matches C_1 for different values of α . Original script: data generated by the original script, $N_R = 10^3$ replications. The frequencies are not statistically different from the ones reported in HP tables 4, 6, 7. Modified script: data from modified script for the generation of AR series, $N_R = 10^4$ replications.

modified script is employed, the regression model is well-specified, i.e. it contains the DGP as a special case.¹²

The original results in HP and those obtained with the modified script in Table 2 document how HP's algorithm depends on α , the significance level chosen in the test R in (5).

3. GSA APPROACH

The HP algorithm uses t -ratios to rank regressors in order of importance, in order to select a subset of model Γ_s . In this study we propose to complement the t -ratios with a GSA measure, called the 'total sensitivity index'. An algorithm is then developed which combines this new ranking with the ranking by t -statistics; we call this the 'GSA algorithm'. Following HP, we define a testing sequence based on this new ranking. Unlike in HP, we adopt a 'bottom-up' selection process which builds candidate models by adding regressors in descending order of importance; this 'bottom-up' selection process has better theoretical properties, see e.g. Paruolo (2001), and can still be interpreted as a GETS procedure. In this section we introduce the total sensitivity index; the description of the GSA algorithm is deferred to Section 4.

The total sensitivity index in GSA is based on systematic MC exploration of the space of the inputs, as is commonly practiced in mathematical modeling in natural sciences and engineering. The 'mechanistic' models in these disciplines are mostly principle-based, possibly involving the solution of some kind of (differential) equation or optimization problem, and the output - being

¹²This finding is similar to the one reported in Hendry and Krolzig (1999), section 6; they re-run HP design using PcGets, and they document similar increases in C_1 catches in DGP 2 and 7 for their modified algorithms. Hence, it is possible that this result is driven by the correction of the script for the generation of the AR series.

the result of a deterministic calculation - does not customarily include an error term. Reviews of global sensitivity analysis methods used therein are given in [Helton et al. \(2006\)](#), [Santner et al. \(2003\)](#), [Saltelli et al. \(2012\)](#).¹³

These techniques are applied here conditionally on the sample $\mathbf{Z} = (\mathbf{y}, \mathbf{X})$ generated as in Section 2.2. Conditionally on the sample \mathbf{Z} , we consider a measure of model fit, such as an information criterion, indicated as $q(\gamma)$. In the application we take $q(\gamma)$ to be BIC.¹⁴ Remark that q is a continuous random variable that depends on the discretely-valued γ (conditionally on \mathbf{Z}).

GSA aims to explore the effect on q when varying inputs γ across the full hyperspace Γ of all possible γ configurations. Specifically, we employ the total sensitivity index of each variable; this can be interpreted as a measure of how much the given regressor contributes to the model fit, as represented by any likelihood-related quantity, such as BIC.

This sensitivity measure belongs to the class of variance-based methods, which aim to decompose the variance of q across models into portions attributable to inputs and sets of inputs. Variance-based measures of sensitivity are the computer-experiment equivalents of the analysis of the variance of in experimental design, see [Archer et al. \(1997\)](#).¹⁵ The objective is to capture both the main effect and the interaction effects of the input factors onto the output q , see [Saltelli et al. \(2012\)](#).

3.1. Sensitivity Measures. Let $\mathcal{Q} = \{q(\gamma), \gamma \in \Gamma\}$ be the set of all possible values of q varying γ in the set of models Γ . Let P be the uniform distribution on $\gamma \in \Gamma$ and \mathbb{P} be the induced probability measure on q ; let also \mathbb{E}, \mathbb{V} indicate the expectation and variance operator with respect to \mathbb{P} . This probability space on the model space is introduced here only to simplify exposition, and it does not correspond to any ex-ante probability on the space of models.

We next partition the γ vector into two components γ_i and γ_{-i} , where γ_{-i} contains all elements in γ except γ_i . We let $\mathbb{E}_a(\cdot|b)$ and $\mathbb{V}_a(\cdot|b)$ (respectively $\mathbb{E}_a(\cdot)$ and $\mathbb{V}_a(\cdot)$) indicate the conditional (respectively marginal) expectation and variance operators with respect to a partition (a, b) of γ , where a and b are taken equal to γ_i and to γ_{-i} .

Variance-based measures rely on decomposing the variance of the output, $V = \mathbb{V}(q)$, into portions attributable to inputs and sets of inputs. There are two commonly-accepted variance-based measures, the ‘first-order sensitivity index’ S_i , [Sobol’ \(1993\)](#), and the ‘total-order sensitivity index’ S_{Ti} , [Homma and Saltelli \(1996\)](#).

The first-order index measures the contribution to V of varying the i -th input alone, and it is defined as $S_i = \mathbb{V}_{\gamma_i}(\mathbb{E}_{\gamma_{-i}}(q | \gamma_i))/V$. This corresponds to seeing the effect of including or not including a regressor, but averaged over all possible combinations of other regressors. This measure does not account for interactions with the inclusion/exclusion of other regressors; hence it is not used in the present paper.

¹³A recent application of these methods to the quality of composite indicators is given in [Paruolo et al. \(2013\)](#).

¹⁴ $q(\gamma)$ is a function of \mathbf{Z} , but we omit to indicate this in the notation.

¹⁵In experimental design the effects of factors are estimated over levels; instead, variance-based methods explore the entire distribution of each factor.

Instead, we focus here on the total effect index, which is defined by [Homma and Saltelli \(1996\)](#) as

$$S_{Ti} = \frac{\mathbb{E}_{\gamma_{-i}} (\mathbb{V}_{\gamma_i} (q \mid \gamma_{-i}))}{V} = 1 - \frac{\mathbb{V}_{\gamma_{-i}} (\mathbb{E}_{\gamma_i} (q \mid \gamma_{-i}))}{V}. \quad (6)$$

In the following we indicate the numerator of S_{Ti} as $\sigma_{Ti}^2 = \mathbb{E}_{\gamma_{-i}} (\mathbb{V}_{\gamma_i} (q \mid \gamma_{-i}))$, and we use the shorthand S_T for S_{Ti} .

Examining σ_{Ti}^2 , one can notice that the inner term, $\mathbb{V}_{\gamma_i} (q \mid \gamma_{-i})$, is the variance of q due inclusion/exclusion of regressor i , but conditional on a given combination γ_{-i} of the remaining regressors. The outer expectation then averages over all values of γ_{-i} ; this quantity is then standardised by V to give the fraction of total output variance caused by the inclusion of x_i . The second expression shows that S_{Ti} is 1 minus the first order effect for γ_{-i} .

These measures are based on the standard variance decomposition formula, or ‘law of total variance’, see e.g. [Billingsley \(1995\)](#), Problem 34.10(b). In the context of GSA these decomposition formulae are discussed in [Archer et al. \(1997\)](#), [Saltelli and Tarantola \(2002\)](#), [Sobol’ \(1993\)](#), [Brell et al. \(2010\)](#). For further reading about GSA in their original setting, we refer to [Saltelli et al. \(2012\)](#).

3.2. Monte Carlo Estimation. In order to calculate the total sensitivity measure S_{Ti} one should be able to compute $q(\gamma)$ for all $\gamma \in \Gamma$, which is unfeasible or undesirable. Instead, S_{Ti} can be estimated by MC, sampling from the space of inputs Γ . Here we select $\gamma_i \in \{0, 1\}$ with $P(\gamma_i = 0) = P(\gamma_i = 1) = 0.5$, independently of γ_{-i} .

This suggests the following MC sampling scheme. Generate a random draw of γ from P , say γ_* ; then consider elements $\gamma_*^{(i)}$ with all elements equal to γ_* except for the i -th coordinate which is switched from 0 to 1 or vice-versa, $\gamma_{*i}^{(i)} = 1 - \gamma_{*i}$. Doing this for each coordinate i generates p additional points $\gamma_*^{(i)}$, and p pairs of γ vectors, γ_* and $\gamma_*^{(i)}$, that differ only in the coordinate i . This is then used to calculate $g(\gamma)$ and apply an ANOVA-like estimation of main effect and residual effects.

More precisely, initialize ℓ at 1; then:

- (1) Generate a draw of γ from P , where γ is a p -length vector with each element is randomly selected from $\{0, 1\}$. Denote this by γ_ℓ .
- (2) Evaluate $q_\ell = q(\gamma_\ell)$.
- (3) Take the i th element of γ_ℓ , and invert it, i.e. set it to 0 if it is 1, and 1 if it is 0. Denote this new vector with inverted i th element as $\gamma_\ell^{(i)}$.
- (4) Evaluate $q_{i\ell} = q(\gamma_\ell^{(i)})$.
- (5) Repeat steps 3 and 4 for $i = 1, 2, \dots, p$.
- (6) Repeat steps 1-5 N times, i.e. for $\ell = 1, 2, \dots, N$.

The computational MC estimator for σ_{Ti}^2 and V are defined as follows, see [Saltelli et al. \(2010\)](#),

$$\hat{\sigma}_{Ti}^2 = \frac{1}{4N} \sum_{\ell=1}^N (q_{i\ell} - q_\ell)^2, \quad \hat{V} = \frac{1}{N-1} \sum_{\ell=1}^N (q_\ell - \bar{q})^2, \quad (7)$$

where $\bar{q} = \frac{1}{N} \sum_{\ell=1}^N q_\ell$. This delivers the following plug-in estimator for S_T , $\hat{S}_{Ti} = \hat{\sigma}_{Ti}^2 / \hat{V}$. Readers familiar with sensitivity analysis may notice that the estimator in (7) is different by a

factor of 2 to the estimator quoted in Saltelli et al. (2010). The reason for this is given in eq. (9) in Appendix A.¹⁶

We investigate the theoretical properties of ordering of variables based on S_T in Appendix A; there we show that these orderings satisfy the following minimal requirement. When the true regressors included in the DGP and the irrelevant ones are uncorrelated, the ordering of regressors based on S_T separates the true from the irrelevant regressors in large samples. One may hence expect this result to apply to other more general situations.¹⁷

We also investigated the contribution of S_T in practice, using HP's experimental designs, as reported in the following section.

3.3. Ordering variables. In order to see how S_T can be used as an alternative or complementary method of ranking regressors, the following numerical experiment was performed. For each of the 11 DGPs under investigation, $N_R = 500$ samples \mathbf{Z} were drawn; on each sample, regressors were ranked by the t -test and S_T , using $N = 128$ in (7). We generically indicate method m for the ordering, where m takes the values t for t -test and S for S_T orderings. Both for the t -test ranking and the S_T ranking, the ordering is from the best-fitting regressor to the worst-fitting one.

In order to measure how successful the two methods were in ranking regressors, we defined the following measure δ of minimum relative covering size. Indicate by $\varphi_0 = \{i_1, \dots, i_{r_0}\}$ the set containing the positions i_j of the true regressors in the list $i = 1, \dots, p$; i.e. for each j one has $\gamma_{0,i_j} = 1$. Recall also that r_0 is the number of elements in φ_0 . Next, for a generic replication j , let $\varphi_\ell^{(m)} = \{i_1^{(m)}, \dots, i_\ell^{(m)}\}$ be the set containing the first ℓ positions $i_j^{(m)}$ induced by the ordering of method m . Let $b_j^{(m)} = \min\{\ell : \varphi_0 \subseteq \varphi_\ell^{(m)}\}$ be the minimum number of elements ℓ for which $\varphi_\ell^{(m)}$ contains the true regressors. We observe that $b_j^{(m)}$ is well defined, because at least for $\ell = p$ one always has $\varphi_0 \subseteq \varphi_p^{(m)} = \{1, \dots, p\}$. We define δ to equal $b_j^{(m)}$ divided by its minimum; this corresponds to the (relative) minimum number of elements in the ordering m that covers the set of true regressors.

Observe that, by construction, one has $r_0 \leq b_j^{(m)} \leq p$, and that ideally one wishes $b_j^{(m)}$ to be as small as possible; ideally one would like to have $b_j^{(m)} = r_0$. Hence for $\delta_j^{(m)}$ defined as $b_j^{(m)}/r_0$ one has $1 \leq \delta_j^{(m)} \leq p/r_0$. We then compute $\delta^{(m)}$ as the average $\delta_j^{(m)}$ over $j = 1, \dots, N_R$, i.e. $\delta^{(m)} = \frac{1}{N_R} \sum_{j=1}^{N_R} \delta_j^{(m)}$.

For example, if the regressors, ranked in descending order of importance by method m in replication j , were $x_3, x_{12}, x_{21}, x_{11}, x_4, x_{31}, \dots$, and the true DGP were x_3, x_{11} the measure δ_j would be 2; in fact the smallest-ranked set containing x_3, x_{11} has 4 elements $b_j^{(m)} = 4$, and $r_0 = 2$.

¹⁶A heuristic reason for this is that the MC method involves a probabilistic exploration of models, and $P(\gamma_i = 0) = P(\gamma_i = 1) = 0.5$. Note that in MC analyses with continuous variables, it is usually advisable to use low-discrepancy sequences due to their space-filling properties, see Sobol' (1967), which give faster convergence with increasing N . However, since γ can only take binary values for each element, low-discrepancy sequences offer no obvious advantage over (pseudo-)random numbers.

¹⁷The results in Appendix A show that one can build examples where the ordering of regressors based on S_T fails to separate the sets of true and irrelevant regressors. The following analysis investigates how often this happens in HP's designs of experiments.

DGP	1	2	3	4	5	6	6A	6B	7	8	9	Mean
S_T		1.00	1.01	1.00	1.00	1.00	1.12	1.02	1.15	1.64	1.13	1.11
t -test		1.00	1.53	1.04	1.00	1.06	3.95	1.14	1.04	1.00	1.01	1.38

TABLE 3. Values of δ for all DGPs (average over 500 data replications per DGP), using t -test and S_T . Mean refers to average across DGPs. Comparatively poor rankings are in boldface.

The results over the $N_R = 500$ replications are summarized in Table 3. Overall S_T appears to perform better than t -ordering. For some DGPs (such as DGP 2 and 5) both approaches perform well ($\delta = 1$ indicating correct ranking for all 500 data sets). There are other DGPs where the performance is significantly different. In particular the t -test is comparatively deficient on DGPs 3 and 10, whereas S_T performs worse on DGP 8. This suggests that there are some DGPs in which S_T may offer an advantage over the t -test in terms of ranking regressors in order of importance. This implies that a hybrid approach, using both measures, may yield a more efficient method of regressor selection, which leads to the model selection algorithm proposed in the following section.

4. GSA ALGORITHM

In this section we present a hybrid approach, which combines the search paths obtained using the t -ratios and the S_T measures, and then selects the best model between the two resulting specifications. The combined procedure is expected to be able to reap the advantages of both orderings. For simplicity, we call this algorithm the GSA algorithm, despite the fact that it embodies some of the characteristics of the HP algorithm. The rest of this section contains a description of the GSA algorithm in its basic form and with two modifications.

4.1. The basic algorithm. The procedure involves ranking the regressors by t -score or S_T , then adopting a ‘bottom up’ approach, where candidate models are built by successively adding regressors in order of importance. The steps are as follows.

- (1) Order all regressors by method m (i.e. either the t -score or S_T).
- (2) Define the initial candidate model as the empty set of regressors.
- (3) Add to the candidate model the highest-ranking regressor (that is not already in the candidate model).
- (4) Perform an F test, comparing the validity of the candidate model to that of the GUM.
- (5) If the p -value of the F test in step 4 is below a given significance level α , go to step 3 (continue adding regressors), otherwise, go to step 6.
- (6) Since the F -test has not rejected the model in step 4, this is the selected model $\gamma^{(m)}$.

In the following, we use the notation $\gamma^{(t)}$ (respectively $\gamma^{(S)}$) when t -ratios (respectively S_T) are used for the ordering. Note that candidate variables are added starting from an empty specification; this is hence a ‘bottom up’ approach.

We observe that this ‘bottom up’ approach is in line with the GETS philosophy of model selection; in fact it corresponds to the nesting of models known the ‘Pantula-principle’ in cointegration rank determination, see [Johansen \(1996\)](#). Every model in the sequence is compared

with the GUM, and hence the sequence of tests can be interpreted as an implementation of the GETS philosophy. Moreover, it can be proved that, for large sample sizes, the sequence selects the smallest true model in the sequence with probability equal to $1 - \alpha$, where α is the size of each test. Letting α tend to 0 as the sample size gets large, one can prove that this delivers a true model with probability tending to 1.¹⁸

As a last step, the final choice of regressors $\hat{\gamma}$ is chosen between $\gamma^{(t)}$ and $\gamma^{(S)}$ as the one with the fewest regressors (since both models have been declared valid by the F -test). If the number of regressors is the same, but the regressors are different, the choice is made using the BIC.

The GSA algorithm depends on some key constants; the significance level of the F -test, α , is a truly ‘sensitive’ parameter, in that varying it strongly affects its performance. Of the remaining constants in the algorithm, N , the number of points in the GSA design, can be increased to improve accuracy; in practice it was found that $N = 128$ provided good results, and further increases made little difference.

4.2. Adaptive- α . Varying α essentially dictates how ‘strong’ the effect of regressors should be to be included in the final model, such that a high α value will tend to include more variables, whereas a low value will cut out variables more harshly. The difficulty is that some DGPs require low α for accurate identification of the true regressors, whereas others require higher values. Hence, there could exist no single value of α that is suitable for the identification of all DGPs.

A proposed modification to deal with this problem is to use an ‘adaptive- α ’, α_ϕ , which is allowed to vary depending on the data. This is based on the observation that the F -test returns a high p -value p_H (typically of the order 0.2-0.6) when the proposed model is a superset of the DGP, but when one or more of the true regressors are missing from the proposed model, the p -value will generally be low, p_L (of the order 10^{-3} say). The values of p_H and p_L will vary depending on the DGP and data set, making it difficult to find a single value of α which will yield good results across all DGPs. For a given DGP and data set, the p_H and p_L values are easy to identify.

Therefore, it is proposed to use a value of α_ϕ , such that for each data set,

$$\alpha_\phi = p_L + \phi(p_H - p_L) \quad (8)$$

where p_H is taken as the p -value resulting from considering a candidate model with all regressors that have $S_{T_i} > 0.01$ against the GUM, and p_L is taken as the p -value from considering the empty set of regressors against the GUM. The reasoning behind the definition of p_H is that it represents a candidate model which will contain the DGP regressors with a high degree of confidence. ϕ is a tuning parameter that essentially determines how far between p_L and p_H the cutoff should be. Figure 1 illustrates this on a data set sampled from DGP 6B. Note that α_ϕ is used in the F -test for both the t -ranked regressors as well as those ordered by S_T .

4.3. Skipping regressors. In order to correct situations where the ordering of the regressors is not the correct one, we present here an extension of the algorithm that allows the possibility to skip regressors in the final model. More precisely, when step 6 is reached, it is allowed to

¹⁸See for instance [Paruolo \(2001\)](#). Recall that any model whose set of regressors contains the true one is ‘true’.




FIGURE 1. p -values from F -test comparing candidate models to the GUM in a sample from DGP 6B, for the 6 highest-ranked regressors. Here $\phi = 0.2$ and α_ϕ is marked as a dotted line.

try removing any of the other remaining regressors, one by one, regardless of the ranking. This approach is used instead of an exhaustive search of the combinations of remaining regressors, because occasionally there may still be too many regressors left to make this feasible.

In Section 5 the performance of the algorithm is examined, with and without the extensions mentioned. We call $S_{T\text{full}}$ the S_T procedure with adaptive- α and skipping regressors; we call $S_{T\text{no-skip}}$ the same procedure without skipping regressors; we call $S_{T\text{simple}}$ the one without adaptive- α and without skipping regressors.

5. RESULTS

In this section we present results, using the performance measures introduced in Section 2.5. The performance is measured with respect to the true DGP or with respect to the effective DGP (EDGP) that one can hope to recover, given the signal to noise ratio. Because the HP and GSA algorithms depend on tunable constants, we give results for various values of these constants.

The procedure employed to define the EDGP is discussed in Appendix B; it implies that the only EDGP differing from the true DGP are DGP 6 and DGP 9. DGP 6 contains regressors 3 and 11, but regressor 3 is weak and EDGP 6 hence contains only regressor 11. DGP 9 contains regressors 3, 11, 21 29 and 37 but regressor 3 and 21 are weak and they are dropped from the corresponding EDGP 9. More details are given in Appendix B.

The HP algorithm depends on the significance levels α , and the GSA algorithm on the threshold ϕ (which controls α_ϕ) for $S_{T\text{no-skip}}$ and $S_{T\text{full}}$ and on α for $S_{T\text{simple}}$. Because the values of α and ϕ can seriously affect the performance of the algorithms, a fair comparison of the performance of the algorithms may be difficult, especially since the true parameter values will not be known in practice. To deal with this problem, the performance of the algorithms was measured at a number of parameter values within a plausible range.

This allowed two ways of comparing the algorithms: first, the ‘optimised’ performance, corresponding to the value of α or ϕ that produced the highest C_1 score, averaged over the 11 DGPs. This can be viewed as the ‘potential performance’. In practice, the optimization was performed with a grid search on α and ϕ with $N_R = 10^3$ replications, averaging across DGPs.

EDGP	$S_{T\text{simple}}$ $\alpha = 0.0371$			$S_{T\text{no-skip}}$ $\phi = 0.3$			$S_{T\text{full}}$ $\phi = 0.3$			$\text{HP}_{\text{optimized}}$ $\alpha = 4 \cdot 10^{-4}$		
	C_1	Gauge	Pot.	C_1	Gauge	Pot.	C_1	Gauge	Pot.	C_1	Gauge	Pot.
1	98.70	0.11	100	99.83	0.01	100	99.83	0.00	100	99.22	0.02	100
2	98.52	0.09	99.98	99.37	0.02	100	99.37	0.02	100	98.94	0.03	100
3	79.36	0.80	94.73	95.23	0.10	98.53	95.97	0.06	98.48	62.01	0.05	81.17
4	98.59	0.09	100	99.16	0.03	100	99.16	0.02	100	99.29	0.02	99.92
5	98.79	0.08	100	99.86	0.00	100	99.86	0.00	100	99.26	0.02	100
6	98.70	0.09	99.99	99.22	0.03	99.99	99.22	0.02	99.99	99.19	0.03	99.84
6A	65.31	0.46	87.86	78.37	0.66	97.92	96.24	0.05	98.52	85.30	0.55	92.91
6B	97.61	0.10	99.99	98.57	0.04	99.99	99.37	0.02	100	98.38	0.07	99.52
7	92.66	0.13	98.58	97.09	0.09	99.87	99.50	0.01	99.90	98.76	0.03	99.82
8	98.44	0.07	99.97	99.91	0.00	100	99.92	0.00	100	99.05	0.03	100
9	91.38	0.18	98.62	96.53	0.11	99.93	99.61	0.01	99.94	98.18	0.04	99.78
Mean	92.55	0.20	98.16	96.65	0.10	99.66	98.91	0.02	99.71	94.33	0.08	97.54

TABLE 4. Percentage C_1 , gauge and potency by EDGP. Optimised parameter values used.

Secondly, a qualitative comparison was drawn between the algorithms of the average performance over the range of parameter values. This latter comparison gives some insight into the more realistic situation where the optimum parameter values are not known.

5.1. Search for the EDGP. Table 4 shows the classification results in terms of C_1 matches, as well as the potency and gauge measures, for both algorithms at their optimal parameter values. Results are shown with and without the extensions discussed in Section 4, using $N_R = 10^4$. Recovery of the true specification is here understood in the EDGP sense.

The C_1 column measures the percentage frequency with which the algorithms identified the EDGP. One notable fact is that the performance of the HP algorithm can be vastly improved (compared to the results in HP) simply by setting α to a better value, in this case $\alpha = 4 \times 10^{-4}$, compare with Table 2.

The comparison shows that with the full S_T algorithm, the correct classification rate, averaged over all DGPs, is improved from 94.6% to 98.9%, which corresponds to five-fold drop in the failure rate from 5.4% to 1.1%.

Removing the ‘skipping’ extension, the performance falls to 97.1%, and further to 92.7% without the adaptive- α feature. Examining the DGPs individually, one can see that the HP algorithm performs well on all DGPs except 3 and 6A, where C_1 is significantly lower than for the GSA algorithm. The S_T method, however, performs well on all the DGPs investigated here, giving improvements over the HP algorithm of around 30 percentage points in DGP3, and about 10 percentage points in DGP 6A. It is evident though that the adaptive- α and the skipping extensions contribute significantly to the performance in those cases.

The potency and gauge measures (also in Table 4) reveal a little more about the nature of the errors made by the algorithms. Gauge is very low for all algorithms, due to the fact that

DGP	$S_{T\text{simple}}$ $\alpha = 0.0371$			$S_{T\text{no-skip}}$ $\phi = 0.3$			$S_{T\text{full}}$ $\phi = 0.3$			HP $\alpha = 4 \cdot 10^{-4}$		
	C_1	Gauge	Pot.	C_1	Gauge	Pot.	C_1	Gauge	Pot.	C_1	Gauge	Pot.
6	1.00	0.09	50.01	1.00	0.03	50.01	2.00	0.02	50.01	3.00	0.03	49.94
9	0.00	0.19	59.20	0.00	0.11	59.97	0.00	0.01	59.97	0.00	0.04	59.92
Mean	0.50	0.14	54.60	0.50	0.07	54.99	1.00	0.02	54.99	1.50	0.03	54.93

TABLE 5. Percentage C_1 , gauge and potency by DGP. Optimised parameter values used. The mean frequency is taken over all DGPs, but only the results for DGPs 6 and 9 are shown since the remaining results are identical to Table 4.

the DGPs consist of only a small fraction of the number of candidate regressors, as well as the good performance of all algorithms. One can see though, that higher gauge measures are found in DGP 6A, indicating the inclusion of irrelevant regressors, except for the full S_T algorithm, which has gauges of practically zero. The potency measures show that the true regressors are being identified nearly all the time, except for the HP algorithm on DGP3, which removes true regressors with some positive frequency.

5.2. Recovering the DGP. Although it is argued here that the signal-to-noise ratio in DGPs 6 and 9 is too low for certain regressors to be identified, it is still worth looking at the results with respect to the true DGP, shown in Table 5. All algorithms failed to identify the true DGP even once out of the 10^4 runs. This fact is reflected in the potency, which drops from 100% to 50% (DGP 6), and about 60% (DGP 9). These results are mirrored in the original results of HP. This suggest that GSA may not help when regressors are ‘weak’.

5.3. Robustness of algorithms. As discussed earlier, the results in Table 4 are obtained after optimisation of the tuning parameters α and ϕ . This provides a measure of potential performance, but in reality the best α and ϕ will not be known. For this reason it is indicative to show the results when varying the tuning parameter. The upper panel in Figure 2 shows how the categorisation of the final model varies with α in the HP algorithm. It is clear that the peak performance of the algorithm is obtained in a small neighborhood around a rather sharp maximum at a low α value – increasing α from this value results in a rapid increase in C_2 , whereas decreasing it sharply increases C_3 .

In contrast, the lower panel in Figure 2 shows the same plot for the full S_T algorithm. While the value of ϕ varies between 0.1 and 0.5, the value of C_1 is generally above 95%, and C_2 and C_3 are consistently very low. While it is difficult to accurately compare this with the HP algorithm, due to the incomparable scales of the two optimising parameters, the S_T algorithm seems to be considerably more robust, and has the advantage that the tuning parameter, ϕ , is relatively problem-independent.

6. CONCLUSIONS

In the model selection problem, one has to choose whether or not to include candidate variables. The approach in this paper is to view the problem as a sensitivity analysis of a measure

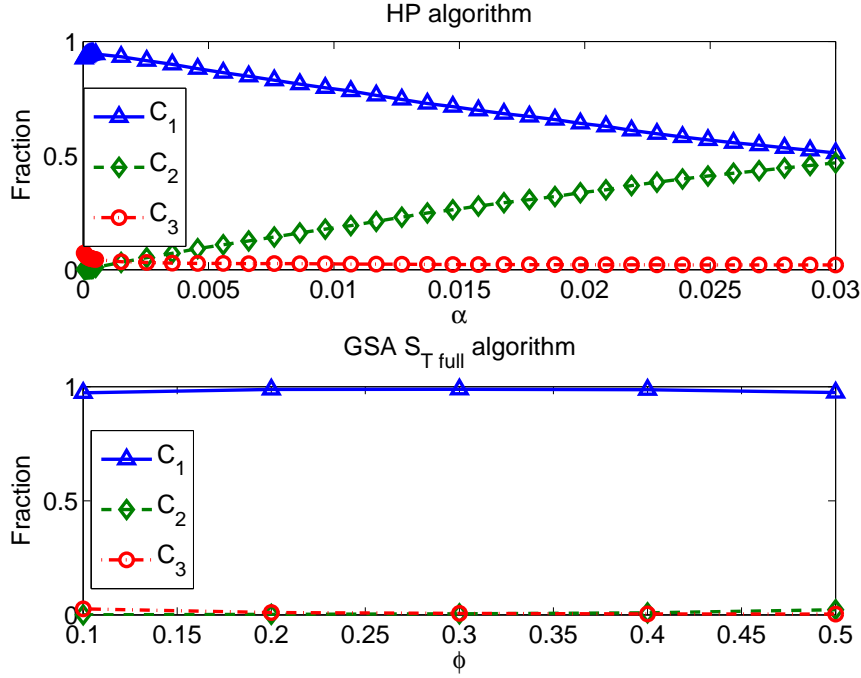


FIGURE 2. Optimisation of algorithms with respect to tuning parameters; upper panel: HP algorithm; lower panel: $S_{T\text{full}}$ algorithm.

of fit on the space of candidate variables. One therefore calculates the sensitivity e.g. of the BIC with respect to the presence (or absence) of each candidate variable. Since global methods are used, the global ‘importance’ is estimated, including all interaction effects among indicators for regressor presence.

These interactions are in principle relevant, as the importance of a given regressor being or not being in the DGP is conditioned by inclusion or exclusion of the other regressors. For this reason we used S_T , a sensitivity measure capable of appreciating the sensitivity of a trigger for the presence of one regressor, inclusive of its interaction effects with triggers for all other regressors.

The GSA algorithm outperforms the HP algorithm both in terms of its potential if tuning parameters were known, and in average performance in the practical situation when tuning parameters are unknown. The improvement is substantial and amounts to a five-fold drop in the failure rate over the ensemble of HP’s designs. Arguably, the robustness of the algorithm is an even more distinguishing feature, since the optimal parameter values would not be known in a practical case.

This study has been a first exploration of new uses of GSA in the world of model selection; it has shown to what extent measures from GSA can contribute, albeit on a small (but representative) set of test problems. It appears then that S_T can be used in selecting important sources of variation in regression. These results call for more research on the use of GSA methods in model selection.

REFERENCES

- Abadir, K. and J. R. Magnus (2002). Notation in econometrics: a proposal for a standard. *The Econometrics Journal* 5, 76–90.
- Anderson, T. W. (1971). *The Statistical Analysis of Time Series*. Wiley.
- Archer, G., A. Saltelli, and I. Sobol (1997). Sensitivity measures, anova-like techniques and the use of bootstrap. *Journal of Statistical Computation and Simulation* 58(2), 99–120.
- Billingsley, P. (1995). *Probability and Measure*. NY: John Wiley & Sons.
- Bittman, R., J. P. Romano, C. Vallarino, and M. Wolf (2009). Testing multiple hypotheses with common effect. *Biometrika* 96, 399–410.
- Box, G. E. P. and N. R. Draper (2007). *Empirical Model Building and Response Surfaces*. Wiley, New York.
- Brell, G., G. Li, and H. Rabitz (2010). An efficient algorithm to accelerate the discovery of complex material formulations. *Journal of Chemical Physics* 132(17), 174103–1–10.
- Brunea, F. (2008). Consistent selection via the Lasso for high dimensional approximating regression models. In B. Clarke and S. Ghosal (Eds.), *Pushing the limits of contemporary statistics: Essays in honor of J. K. Gosh*, pp. 122–137. Dordrecht: IMS.
- Burnham, K. P. and D. R. Anderson (2002). *Model Selection and Multimodel Inference - A Practical Information-theoretic Approach*, 2nd ed. Springer.
- Castle, J. L., J. A. Doornik, and D. F. Hendry (2011). Evaluating automatic model selection. *Journal of Time Series Econometrics* 3.
- Claeskens, G. and N. L. Hjort (2003). The focused information criterion, with discussion. *Journal of the American Statistical Association* 98, 900–945.
- Danilov, D. and J. Magnus (2004). On the harm that ignoring pretesting can cause. *Journal of Econometrics* 122, 27–46.
- Foster, D. P. and E. I. George (1994). The risk inflation criterion for multiple regression. *Annals of Statistics* 22, 1947–1975.
- Freedman, D. (1983). A note on screening regression equations. *The American Statistician* 37, 152–155.
- Freedman, D. and P. Humphreys (1999). Are there algorithms that discover causal structure? *Synthese* 121, 29–54.
- Freedman, L. E. and D. Pee (1989). Return to a note on screening regression equations. *The American Statistician* 43, 279–282.
- Freedman, L. E., D. Pee, and N. Midthune (1992). The problem of underestimating the residual error variance in forward stepwise regression. *Journal of the Royal Statistical Society, Series D (The Statistician)* 41, 405–412.
- Hansen, B. (1999). Discussion of ‘Data Mining Reconsidered’ by K.D. Hoover and S.J. Perez. *Econometrics Journal* 2, 192–201.
- Harrell, F. (2001). *Regression Modeling Strategies*. Springer.
- Helton, J., J. Johnson, C. Sallaberry, and C. Storlie (2006). Survey of sampling-based methods for uncertainty and sensitivity analysis. *Reliability Engineering and System Safety* 91 (10–11), 1175–1209.

- Hendry, D. and H.-M. Krolzig (1999). Improving on ‘Data Mining Reconsidered’ by K.D. Hoover and S.J. Perez. *Econometrics Journal* 2, 41–58.
- Hendry, D. F. and A. Krolzig (2005). The properties of automatic gets modelling. *Economic Journal* 115, C32–C61.
- Hjort, N. L. and G. Claeskens (2003). Frequentist model average estimators. *Journal of the American Statistical Association* 98, 879–899.
- Homma, T. and A. Saltelli (1996). Importance measures in global sensitivity analysis of nonlinear models. *Reliability Engineering & System Safety* 52(1), 1–17.
- Hoover, K. and S. Perez (1999). Data mining reconsidered: encompassing and the general-to-specific approach to specification search. *The Econometrics Journal* 2(2), 167–191.
- Imbens, G. and J. Wooldridge (2009). Recent developments in the econometrics of program evaluation. *Journal of Economic Literature* 47, 5–86.
- Johansen, S. (1996). *Likelihood-based Inference in Cointegrated Vector Auto-Regressive Models*. Oxford University Press.
- Leamer, E. E. (1983). Let’s take the con out of econometrics. *The American Economic Review* 73(1), pp. 31–43.
- Leeb, H. and B. M. Poetscher (2006). Can one estimate the conditional distribution of post-model-selection estimators? *Annals of Statistics* 34, 2554–2591.
- Liu, W. and Y. Yang (2011). Parametric or nonparametric? A parametricness index for model selection. *The Annals of Statistics* 39(4), 2074–2102.
- Lovell, M. C. (1983). Data mining. *The Review of Economics and Statistics* 65, 1–12.
- Magnus, J. (2007). *Boumans, M.J. (ed.) Measurement in Economics*, Chapter Local sensitivity in econometrics, pp. 295–319. Academic Press: San Diego.
- Magnus, J. and J. Durbin (1999). Estimation of regression coefficients of interest when other regression coefficients are of no interest. *Econometrica* 67, 639–643.
- Magnus, J., O. Powell, and P. Prufer (2010). A comparison of two model averaging techniques with an application to growth empirics. *Journal of Econometrics* 154, 139–153.
- Magnus, J. and A. Vasnev (2007). Local sensitivity and diagnostic tests. *The Econometrics Journal* 10, 166–192.
- Miller, A. (2002). *Subset selection in regression, 2nd ed.* Chapman and Hall, CRC Press.
- Paruolo, P. (2001). The power of lambda max. *Oxford Bulletin of Economics and Statistics* 63, 395–403.
- Paruolo, P. and A. Rahbek (1999). Weak exogeneity in I(2) VAR systems. *Journal of Econometrics* 93, 281–308.
- Paruolo, P., A. Saltelli, and M. Saisana (2013). Ratings and rankings: Voodoo or science? *Journal of the Royal Statistical Society, Series A (Statistics in Society)* 176, 609–634.
- Paulsen, J. (1984). Order determination of multivariate autoregressive time series with unit roots. *Journal of Time Series Analysis* 5, 115–127.
- Phillips, P. C. B. (1997). Econometric model determination. *Econometrica* 64, 763–812.
- Phillips, P. C. B. (2003). Laws and limits of econometrics. *Economic Journal* 113, C26–C52.
- Poetscher, B. M. (1991). Effects of model selection on inference. *Econometric Theory* 7, 163–185.

- Romano, J. P. and M. Wolf (2005). Stepwise multiple testing as formalized data snooping. *Econometrica* 73, 1237–1282.
- Sala-i-Martin, X. (1997). I just ran two million regressions. *The American Economic Review, Papers and Proceedings of the Hundred and Fourth Annual Meeting of the American Economic Association* 87, 178–183.
- Saltelli, A. (2002). Sensitivity analysis for importance assessment. *Risk Analysis* 22, 579–590.
- Saltelli, A., T. Andres, and T. Homma (1993). Sensitivity analysis of model output: An investigation of new techniques. *Computational Statistics & Data Analysis* 15, 211–238.
- Saltelli, A., P. Annoni, I. Azzini, F. Campolongo, M. Ratto, and S. Tarantola (2010). Variance based sensitivity analysis of model output. design and estimator for the total sensitivity index. *Computer Physics Communications* 181 (2), 259–270.
- Saltelli, A., M. Ratto, S. Tarantola, and F. Campolongo (2012). Sensitivity analysis for chemical models. *Chemical Reviews* 112, PR1–PR21.
- Saltelli, A. and S. Tarantola (2002). On the relative importance of input factors in mathematical models. *Journal of the American Statistical Association* 97(459), 702–709.
- Santner, T., B. Williams, and W. Notz (2003). *Design and Analysis of Computer Experiments*. Springer-Verlag.
- Sobol', I. M. (1967). On the distribution of points in a cube and the approximate evaluation of integrals. *USSR Computational Mathematics and Mathematical Physics* 7(4), 86–112.
- Sobol', I. M. (1993). Sensitivity estimates for nonlinear mathematical models. *Mathematical Modeling and Computational Experiment* 1 (4), 407–414.
- Young, P., S. Parkinson, and M. Lees (1996). Simplicity out of complexity: Occam's razor revisited. *Journal of Applied Statistics* 23, 165–210.

APPENDIX A: PROPERTIES OF ORDERINGS BASED ON S_{Ti}

Recall that $S_{Ti} = \sigma_{Ti}^2/V = \mathbb{E}_{\gamma_{-i}}(\mathbb{V}_{\gamma_i}(q | \gamma_{-i}))/V$. In this Appendix, we first express σ_{Ti}^2 as a sum of terms involving $\hat{\sigma}_{\gamma}^2$ for $\gamma \in \Gamma$ in Lemma 1; next we show the large n behavior of $\hat{\sigma}_{\gamma}^2$ in Lemma 2. Lemma 3 states the probability limit of σ_{Ti}^2 . Lemma 4 shows that, in case the true regressors in the DGP and the irrelevant ones are uncorrelated, $\sigma_{Ti}^2 \xrightarrow{p} 0$ for an irrelevant regressor i , while $\sigma_{Ti}^2 \xrightarrow{p} c_i > 0$ for a relevant one. Under the same conditions of Lemma 4, Theorem 5 shows that for large samples, a scree plot on the ordered S_{Ti} allows to separate the relevant regressors from the irrelevant ones.

Let $\gamma_i = \mathbf{e}_i' \gamma$ and $\gamma_{-i} = \mathbf{A}_i' \gamma$, where \mathbf{e}_i is the i -th column of the identity matrix of order p , \mathbf{I}_p and \mathbf{A}_i is a $p \times p - 1$ matrix containing all the columns of \mathbf{I}_p except the i -th one. We write $q(\gamma)$ as $q(\gamma_i, \gamma_{-i})$ or, more simply as $q_{-i}(\gamma_i)$. Denote by $\gamma^{(i,0)}$ the vector corresponding to $\gamma_i = 0$, with the remaining coordinates equal to γ_{-i} , and let $\gamma^{(i,1)}$ the vector corresponding to $\gamma_i = 1$ with the remaining coordinates equal to γ_{-i} . Finally let $\Gamma_{-i} := \{\gamma_{-i} = \mathbf{A}_i' \gamma, \gamma \in \Gamma\}$.

Lemma 1 (σ_{Ti}^2 as an average over γ_i). *One has*

$$\mathbb{E}_{\gamma_{-i}}(\mathbb{V}_{\gamma_i}(q | \gamma_{-i})) = \frac{1}{4 \cdot 2^{p-1}} \sum_{\gamma_{-i} \in \Gamma_{-i}} (q_{-i}(1) - q_{-i}(0))^2 \quad (9)$$

and for q equal to BIC (or any other consistent information criterion)

$$q_{-i}(1) - q_{-i}(0) = \log \left(\frac{\hat{\sigma}_{\gamma^{(i,1)}}^2}{\hat{\sigma}_{\gamma^{(i,0)}}^2} \right) + o(1), \quad (10)$$

where $o(1)$ is a term tending to 0 for large n and $\hat{\sigma}_{\gamma}^2 := n^{-1} \hat{\varepsilon}_{\gamma}' \hat{\varepsilon}_{\gamma}$ where $\hat{\varepsilon}_{\gamma}$ are the residuals of model γ .

Proof. Note that for $h = 1, 2$ one has $\mathbb{E}_{\gamma_i}(q^h | \gamma_{-i}) = \frac{1}{2}(q_{-i}^h(1) + q_{-i}^h(0))$ so that

$$\begin{aligned} \mathbb{V}_{\gamma_i}(q | \gamma_{-i}) &= \mathbb{E}_{\gamma_i}(q^2 | \gamma_{-i}) - (\mathbb{E}_{\gamma_i}(q | \gamma_{-i}))^2 \\ &= \frac{1}{2}(q_{-i}^2(1) + q_{-i}^2(0)) - \frac{1}{4}(q_{-i}^2(1) + q_{-i}^2(0) + 2q_{-i}(1)q_{-i}(0)) \\ &= \frac{1}{4}(q_{-i}(1) - q_{-i}(0))^2. \end{aligned}$$

Hence one finds (9). When q is BIC, $q(\gamma) = \log \hat{\sigma}_{\gamma}^2 + k_{\gamma} c_n$ with $c_n := \log(n)/n$. Other consistent information criteria replace $\log n$ with some other increasing function $f(n)$ of n with the property $c_n = f(n)/n \rightarrow 0$, see Paulsen (1984) Theorem 1. Note also that $k_{\gamma^{(i,1)}} - k_{\gamma^{(i,0)}} = 1$, and that one has

$$q_{-i}(1) - q_{-i}(0) = \log \left(\frac{\hat{\sigma}_{\gamma^{(i,1)}}^2}{\hat{\sigma}_{\gamma^{(i,0)}}^2} \right) + (k_{\gamma^{(i,1)}} - k_{\gamma^{(i,0)}}) c_n = \log \left(\frac{\hat{\sigma}_{\gamma^{(i,1)}}^2}{\hat{\sigma}_{\gamma^{(i,0)}}^2} \right) + c_n.$$

Because $c_n \rightarrow 0$, one finds (10). \square

We next wish to discuss the asymptotic behaviour of $\hat{\sigma}_{\gamma}^2$. Let $w_t := (y_t, x_{1,t}, \dots, x_{p,t}, \varepsilon_t)'$, where, without loss of generality, we assume that all variables have mean zero. Denote $\Sigma := E(w_t w_t')$, where

$$\Sigma = \begin{pmatrix} \Sigma_{yy} & \Sigma_{yx} & \sigma^2 \\ \Sigma_{yx} & \Sigma_{xx} & \mathbf{0} \\ \sigma^2 & \mathbf{0} & \sigma^2 \end{pmatrix} = \begin{pmatrix} \Sigma_{yy} & \Sigma_{y1} & \dots & \Sigma_{yp} & \sigma^2 \\ & \Sigma_{11} & & \Sigma_{1p} & 0 \\ & & \ddots & & \\ & & & \Sigma_{pp} & 0 \\ & & & & \sigma^2 \end{pmatrix}.$$

Let $\Sigma_{ij,v} := \Sigma_{ij} - \Sigma_{iv} \Sigma_{vv}^{-1} \Sigma_{vj}$ indicate partial covariances, where $v := \{i_1, \dots, i_s\}$ indicates a set of indices. Note that $\Sigma_{x\varepsilon} = \mathbf{0}$.

Let $\mathbb{J} := \{1, \dots, p\}$ be the set of the first p integers, $\mathbb{T} := \{i \in \mathbb{J} : \beta_{0,i} \neq 0\}$, the set of all regressor indices in the DGP, with r_0 elements, and $\mathbb{M} := \mathbb{J} \setminus \mathbb{T}$ the set of all regressor indices for irrelevant regressors.¹⁹ For each γ , let $a_{\gamma} := \{i_1, \dots, i_{k_{\gamma}}\}'$ indicate the set of indices i_j such that $\gamma_{i_j} = 1$ in γ . Similarly let $b_{\gamma} := \{i_1, \dots, i_s\}'$ indicate the set of indices i_j that belong to $a_{\gamma} \setminus \mathbb{T}$.

We represent β_0 as $\beta_0 = \mathbf{H}\phi$, where \mathbf{H} contains the r_0 columns of \mathbf{I}_p corresponding to $\beta_{0,i} \neq 0$, and ϕ contains the corresponding $\beta_{0,i}$ coefficients. Moreover we write the matrix of regressors in the γ specification as $\mathbf{X} \mathbf{U}_{\gamma}$, where \mathbf{U}_{γ} contains the columns of \mathbf{I}_p with column indices a_{γ} . Define also $\mathbf{M}_{\gamma} := \mathbf{I}_n - \mathbf{X} \mathbf{U}_{\gamma} (\mathbf{U}_{\gamma}' \mathbf{X}' \mathbf{X} \mathbf{U}_{\gamma})^{-1} \mathbf{U}_{\gamma}' \mathbf{X}'$.

¹⁹Here $\mathbb{J} \setminus \mathbb{T}$ denotes the set difference $\mathbb{J} \setminus \mathbb{T} := \{i : i \in \mathbb{J}, i \notin \mathbb{T}\}$; sums over empty sets are understood to be equal to 0.

Lemma 2 (Large sample behavior of $\hat{\sigma}_\gamma^2$). *As $n \rightarrow \infty$, one has*

$$\hat{\sigma}_\gamma^2 \xrightarrow{p} \sigma^2 + \sum_{h,j \in \mathbb{T} \setminus a_\gamma} \beta_{0,h} \Sigma_{hj} b_{\gamma,j}, \quad (11)$$

where $\mathbb{T} \setminus a_\gamma$ is the set of indices of the true regressors omitted from the γ specification, and b_γ is the set of indices $a_\gamma \setminus \mathbb{T}$ of the regressors included in the γ specification except the ones that belong to the DGP. Remark that the sum in (11) is equal to 0 when γ is correctly specified (i.e. it contains all regressors in the DGP) i.e. $\mathbb{T} \setminus a_\gamma = \emptyset$.

Proof. Because $\mathbf{y} = \mathbf{X}\mathbf{H}\phi + \varepsilon$ one has

$$\hat{\sigma}_\gamma^2 = n^{-1} \mathbf{y}' \mathbf{M}_\gamma \mathbf{y} = n^{-1} \varepsilon' \mathbf{M}_\gamma \varepsilon + 2n^{-1} \varepsilon' \mathbf{M}_\gamma \mathbf{X} \mathbf{H} \phi + n^{-1} \phi' \mathbf{H}' \mathbf{X}' \mathbf{M}_\gamma \mathbf{X} \mathbf{H} \phi$$

Because $\Sigma_{x\varepsilon} = \mathbf{0}$, by the law of large numbers for stationary linear processes, see [Anderson \(1971\)](#), one finds

$$\begin{aligned} n^{-1} \varepsilon' \mathbf{M}_\gamma \varepsilon &\xrightarrow{p} \sigma^2 - \Sigma_{\varepsilon x} \mathbf{U}_\gamma (\mathbf{U}_\gamma' \Sigma_{xx} \mathbf{U}_\gamma)^{-1} \mathbf{U}_\gamma' \Sigma_{x\varepsilon} = \sigma^2, \\ n^{-1} \varepsilon' \mathbf{M}_\gamma \mathbf{X} &\xrightarrow{p} \Sigma_{\varepsilon x} (\mathbf{I}_p - \mathbf{U}_\gamma (\mathbf{U}_\gamma' \Sigma_{xx} \mathbf{U}_\gamma)^{-1} \mathbf{U}_\gamma' \Sigma_{xx}) = \mathbf{0}. \end{aligned}$$

Similarly

$$\begin{aligned} n^{-1} \mathbf{H}' \mathbf{X}' \mathbf{M}_\gamma \mathbf{X} \mathbf{H} &\xrightarrow{p} \mathbf{H}' (\Sigma_{xx} - \Sigma_{xx} \mathbf{U}_\gamma (\mathbf{U}_\gamma' \Sigma_{xx} \mathbf{U}_\gamma)^{-1} \mathbf{U}_\gamma' \Sigma_{xx}) \mathbf{H} \\ &= \mathbf{H}' \mathbf{V}_\gamma (\mathbf{V}_\gamma' \Sigma_{xx}^{-1} \mathbf{V}_\gamma)^{-1} \mathbf{V}_\gamma' \mathbf{H} \end{aligned}$$

where $\mathbf{V}_\gamma = \mathbf{U}_{\gamma,\perp}$ contains the columns in \mathbf{I}_p not contained in \mathbf{U}_γ , and the last equality is a special case of a non-orthogonal projection identity, see e.g. eq. (2.13) in [Paruolo and Rahbek \(1999\)](#) and references therein. Here \mathbf{U}_\perp indicates a basis of the orthogonal complement of the space spanned by the columns in \mathbf{U} . Observe that the $(p - k_\gamma) \times r_0$ matrix $\mathbf{C}_\gamma := \mathbf{V}_\gamma' \mathbf{H}$ contains the columns of \mathbf{I}_{p-r_γ} corresponding to the index set of regressors in $v_\gamma := \mathbb{T} \setminus a_\gamma$. Hence, using e.g. eq. (A.4) in [Paruolo and Rahbek \(1999\)](#), one finds $(\mathbf{V}_\gamma' \Sigma_{xx}^{-1} \mathbf{V}_\gamma)^{-1} = \Sigma_{v_\gamma v_\gamma} b_{\gamma,b_\gamma}$. Substituting one finds

$$n^{-1} \phi' \mathbf{H}' \mathbf{X}' \mathbf{M}_\gamma \mathbf{X} \mathbf{H} \phi \xrightarrow{p} \phi' \mathbf{C}_\gamma' \Sigma_{v_\gamma v_\gamma} b_{\gamma,b_\gamma} \mathbf{C}_\gamma \phi.$$

Simplifying one obtains (11). □

The above results lead to the following general formulation of the probability limit of σ_{Ti}^2 .

Lemma 3 (Large sample behaviour of σ_{Ti}^2). *As $n \rightarrow \infty$ one has*

$$\sigma_{Ti}^2 \xrightarrow{p} \frac{1}{4 \cdot 2^{p-1}} \sum_{\gamma_{-i} \in \Gamma_{-i}} \log \left(\frac{\sigma^2 + \sum_{h,j \in \mathbb{T} \setminus a_{\gamma(i,1)}} \beta_{0,h} \Sigma_{hj} b_{\gamma(i,1),j} \beta_{0,j}}{\sigma^2 + \sum_{h,j \in \mathbb{T} \setminus a_{\gamma(i,0)}} \beta_{0,h} \Sigma_{hj} b_{\gamma(i,0),j} \beta_{0,j}} \right)$$

where a_γ is the set of indices of the regressors in the γ specification, and $b_\gamma := a_\gamma \setminus \mathbb{T}$ includes the indices of regressors included in the γ specification except the ones that belong to the DGP.

Proof. Apply Lemma 1 and 2. □

Lemma 3 shows that the limit behavior of σ_{Ti}^2 depends on the covariance structure Σ . Some covariance structures imply that, in the limit, the value of S_T for true regressors is greater than the value of S_T for irrelevant regressors. There also exist other covariance structures which

can imply a reverse ordering.²⁰ In the special case when true and irrelevant regressors are uncorrelated, the next Lemma 4 shows that S_T converges to 0 for irrelevant regressors, while S_T converges to a positive constant for true regressors. This result is then used in Theorem 5 to show that the ordering based on S_T separates true and irrelevant regressors in this special case.

Lemma 4 (Orthogonal regressors in \mathbb{M} and \mathbb{T}). *Assume that $\Sigma_{\ell j} = 0$ for all $j \in \mathbb{T}$ and $\ell \in \mathbb{M}$. Then when $i \in \mathbb{M}$ one has, as $n \rightarrow \infty$, $\sigma_{T_i}^2 \xrightarrow{p} 0$, whereas otherwise when $i \in \mathbb{T}$ one finds*

$$\sigma_{T_i}^2 \xrightarrow{p} c_i > 0. \quad (12)$$

Proof. From Lemma 2, one finds

$$q_{-i}(1) - q_{-i}(0) = \log \left(\frac{\sigma^2 + \sum_{h,j \in \mathbb{T} \setminus a_{\gamma(i,1)}} \beta_{0,h} \Sigma_{hj} b_{\gamma(i,1)} \beta_{0,j}}{\sigma^2 + \sum_{h,j \in \mathbb{T} \setminus a_{\gamma(i,0)}} \beta_{0,h} \Sigma_{hj} b_{\gamma(i,0)} \beta_{0,j}} \right) + o_p(1), \quad (13)$$

Assume that $\Sigma_{\ell j} \neq 0$ for some $j \in \mathbb{T}$ and $\ell \in \mathbb{M}$; then for some $\gamma_{-i} \in \Gamma_{-i}$ one has $\Sigma_{hj} b_{\gamma(i,1)} \neq \Sigma_{hj} b_{\gamma(i,0)}$ in the numerator and denominator on the r.h.s. of (13); let $c \neq 1$ indicate the corresponding ratio. Hence $(q_{-i}(1) - q_{-i}(0))^2$ converges in probability to $\log^2 c > 0$, and because the terms in $\mathbb{E}_{\gamma_{-i}}(\mathbb{V}_{\gamma_i}(q \mid \gamma_{-i})) = \frac{1}{4 \cdot 2^{p-1}} \sum_{\gamma_{-i} \in \Gamma_{-i}} (q_{-i}(1) - q_{-i}(0))^2$, see Lemma 3, are non-negative, one concludes that $\sigma_{T_i}^2 \xrightarrow{p} c_i > 0$.

Assume instead that $\Sigma_{\ell j} = 0$ for all $j \in \mathbb{T}$ and $\ell \in \mathbb{M}$ and $i \in \mathbb{M}$. Then $\mathbb{T} \setminus a_{\gamma(i,\cdot)} = \mathbb{T} \setminus a_{\gamma_{-i}}$ and, because $\Sigma_{\ell j} = 0$ for all $j \in \mathbb{T}$ and $\ell \in \mathbb{M}$, one has $\Sigma_{jb_{\gamma(i,\cdot)}} = 0$. This implies $\Sigma_{hj} b_{\gamma(i,\cdot)} := \Sigma_{hj} - \Sigma_{hb_{\gamma(i,\cdot)}} \Sigma_{b_{\gamma(i,\cdot)} b_{\gamma(i,\cdot)}}^{-1} \Sigma_{b_{\gamma(i,\cdot)} j} = \Sigma_{hj}$. Hence

$$q_{-i}(1) - q_{-i}(0) = \log \left(\frac{\sigma^2 + \sum_{h,j \in \mathbb{T} \setminus a_{\gamma_{-i}}} \beta_{0,h} \Sigma_{hj} \beta_{0,j}}{\sigma^2 + \sum_{h,j \in \mathbb{T} \setminus a_{\gamma_{-i}}} \beta_{0,h} \Sigma_{hj} \beta_{0,j}} \right) + o_p(1) = o_p(1),$$

for all $\gamma_{-i} \in \Gamma_{-i}$ because the numerator and denominator are identical. Thus $(q_{-i}(1) - q_{-i}(0))^2$ converges in probability to $\log^2 1 = 0$ for all $\gamma_{-i} \in \Gamma_{-i}$, and this implies $\sigma_{T_i}^2 \xrightarrow{p} 0$. \square

The following theorem shows that for large samples, a scree plot on the ordered S_{T_i} allows to separate the relevant regressors from the irrelevant ones when true and irrelevant regressors are uncorrelated.

Theorem 5 (Ordering based on S_{T_i} works for orthogonal regressors in \mathbb{M} and \mathbb{T}). *Assume $\Sigma_{\ell j} = 0$ for all $j \in \mathbb{T}$ and $\ell \in \mathbb{M}$ as in Lemma 4. Define $(S_{T(1)}, S_{T(2)}, \dots, S_{T(p)})$ as the set of S_{T_i} values in decreasing order, with $S_{T(1)} \geq S_{T(2)} \geq \dots \geq S_{T(p)}$. Then as $n \rightarrow \infty$ one has*

$$(S_{T(1)}, S_{T(2)}, \dots, S_{T(p)}) \xrightarrow{p} (c_{(1)}, c_{(2)}, \dots, c_{(r_0)}, 0, \dots, 0)$$

where $(c_{(1)}, c_{(2)}, \dots, c_{(r_0)})$ is the set of c_i values defined in (12) in decreasing order. Hence the ordered S_{T_i} values separate the block of true regressors in \mathbb{T} in the first r_0 positions and the irrelevant ones in the last $p - r_0$ ones.

Proof. Direct application of Lemma 4. \square

²⁰Worked out examples illustrating both situations are available from the authors upon request.

APPENDIX B: EFFECTIVE DGP

In this appendix we describe how the notion of ‘weak regressors’ was made operational in the present context. We employ a recently-introduced measure known as the ‘Parametricness Index’ (PI), [Liu and Yang \(2011\)](#), to identify the ‘effective DGP’ (EDGP). Parametricness, in the sense of Liu and Yang, is a measure dependent both on sample size *and* the proposed model; a model is parametric if omission of any of its variables implies a marked change in its fit, and nonparametric otherwise.²¹ Here we take parametricness as a sign of detectability, i.e. of a sufficiently high signal-noise ratio. We apply this concept both to complete specifications as well as to single regressors; in particular we define the EDGP as the subset of DGP regressors which the PI would classify as parametric. Details are given in the following.

Considering a model $\gamma_k \in \Gamma$, one can express the regression fit as $\hat{\mathbf{y}}_k = \mathbf{P}_k \mathbf{y}$, where \mathbf{P}_k is the projection matrix on $\text{col}(\mathbf{X}\mathbf{D}_{\gamma_k})$, and col indicates the column space; let r_{γ_k} be the dimension of $\text{col}(\mathbf{X}\mathbf{D}_{\gamma_k})$. The index PI is defined in terms of an information criterion IC , which depends on λ_n , d and $\hat{\sigma}^2$. Here λ_n is a nonnegative sequence that satisfies $\lambda_n \geq (\log n)^{-1}$, d is a nonnegative constant and $\hat{\sigma}^2$ is a consistent estimator of σ^2 such as $\|\mathbf{y} - \hat{\mathbf{y}}_k\|^2 / (n - r_{\gamma_k})$ with γ_k consistent for γ_0 . In our application we used $\gamma_k = \gamma_0$. The information criterion IC is defined by

$$IC_{\lambda_n, d}(\gamma_k, \hat{\sigma}^2) = \|\mathbf{y} - \hat{\mathbf{y}}_k\|^2 + \lambda_n \log(n) r_k \hat{\sigma}^2 - n \hat{\sigma}^2 + d n^{1/2} \log(n) \hat{\sigma}^2 \quad (14)$$

where $\|\cdot\|$ represents Euclidean distance; here we take $\lambda_n = 1$ and $d = 0$ as suggested in [Liu and Yang \(2011\)](#).

Let now γ_0 be the DGP; PI is now defined in the present context as,

$$PI = \begin{cases} \inf_{\gamma_k \in \Gamma_1(\gamma_0)} \frac{IC_{\lambda_n, d}(\gamma_k, \hat{\sigma}^2)}{IC_{\lambda_n, d}(\gamma_0, \hat{\sigma}^2)} & \text{if } r_{\gamma_0} > 1 \\ n & \text{if } r_{\gamma_0} = 1 \end{cases} \quad (15)$$

where $\Gamma_1(\gamma_0)$ is the set of submodels γ_k of the DGP γ_0 such that $r_{\gamma_k} = r_{\gamma_0} - 1$, i.e. all submodels obtained by removing one regressor at a time (with replacement).²²

The reasoning is that if the model is parametric (and correctly specified for the data), removing any of the regressors will have a marked impact on IC . In contrast, if (some of the) regressors are just incremental terms in a nonparametric approximation, removing one of these regressors will have little effect on IC . [Liu and Yang \(2011\)](#) show that PI converges to 1 for a nonparametric scenario, and goes to infinity in a parametric scenario. The authors suggest to take $PI = 1.2$ is a cutoff point between parametric and nonparametric scenarios; we adopt this threshold in the following.

As suggested by Liu and Yang, PI can, “given the regression function and the noise level ... indicate whether the problem is practically parametric/nonparametric at the current sample size”. If the PI value is close to or below 1, one could conclude that at least some of the terms are ‘undetectable’ at the given sample size, therefore it may be unreasonable to expect an algorithm to identify the DGP correctly.

²¹For example, consider a data set generated by a sine function, with added noise. If it is proposed to model this with a quadratic equation, the data/model should be considered non-parametric. However, if the proposed model included sinusoidal terms, it should be considered parametric.

²²In the original paper γ_0 is replaced by the model $\hat{\gamma}_k$ selected by a weakly consistent information criterion, such as BIC.

DGP	DGP Indices	$F_N(1.2)$	$PI_{0.01}$	$PI_{0.1}$	$E_N(PI)$	$PI_{0.9}$	$PI_{0.99}$	EDGP Indices
1	{}	-	-	-	-	-	-	{}
2	{37}	0.00	16.55	25.59	41.80	60.83	84.61	{37}
3	{37,38}	0.04	0.88	1.53	2.54	3.52	4.17	{37,38}
4	{11}	0.00	30.50	37.82	49.19	61.78	74.88	{11}
5	{3}	0.00	365.84	415.39	493.63	578.17	668.84	{3}
6	{3,11}	0.98	0.37	0.38	0.53	0.79	1.40	{11}
6A	{3,11}	0.00	2.77	4.15	6.44	8.95	11.69	{3,11}
6B	{3,11}	0.00	15.11	18.10	23.04	28.38	33.72	{3,11}
7	{11,29,37}	0.00	2.84	4.16	6.46	8.96	11.76	{11,29,37}
8	{3,21,37}	0.00	5.77	8.40	13.49	19.22	26.41	{3,21,37}
9	{3,11,21,29,37}	1.00	0.75	0.75	0.77	0.81	0.93	{11,29,37}

TABLE 6. Distribution of PI values for DGPs 1-9. $F_n(\cdot)$ is the MC cumulative distribution function of PI and PI_α is the α -quantile of $F_m(\cdot)$. DGPs where EDGP \neq DGP are in boldface.

We apply PI at the level of each DGP; if PI indicates that the DGP is nonparametric, we also investigate which of the submodel is responsible for this and label the corresponding omitted variables as ‘weak’. As in the rest of the paper, we employ a MC approach. We generated 5000 datasets from each DGP and calculated PI for each sample, hence obtaining a distribution of PI values. Table 6 summarizes the MC distribution of PI values through the empirical distribution function $F_m(x) = m^{-1} \sum_{j=1}^m 1(PI_j \leq x)$, where $m = N_R$ and PI_j is the PI value in replication $j = 1, \dots, N_R$. Quantiles of PI are indicated as PI_α , with $\alpha = 0.01, 0.1, 0.9, 0.99$, and the MC mean PI is indicated as $E_N(PI)$, where for simplicity we drop the subscript R from N_R .

The reference threshold is $PI = 1.2$, and $F_N(1.2)$ shows the frequency of PI being below this limit; in other words this gives an estimate for the DGP to be classified as nonparametric. There is a very clear distinction: DGPs 6 and 9 are regarded as nonparametric 98% and 100% of the time respectively. In contrast, all other DGPs are always regarded as parametric, with the slight exception of DGP3, which is a little less clear cut.

Examining the quantiles, DGP3 has a mean PI value of 2.54 and $PI_{0.1} = 1.53$, which puts it in the parametric class in the large majority of cases. DGP 6 has a mean PI of 0.53, and $PI_{0.9} = 0.79$, making it almost always nonparametric. DGP 9 has $PI_{0.99} = 0.93$, making it the most obviously nonparametric DGP. Of the remaining DGPs, all are well above the threshold and can be safely considered parametric.

We next further investigate which regressors are causing the nonparametricity, i.e. which regressors are ‘weak’. We examine the individual IC ratios for each regressor of a given DGP, see (15). Here we let $ICR(i)$ indicate the IC ratio between the DGP and the submodel of the DGP where variable i is removed. Table 7 reports the distribution of $ICR(i)$ for DGPs 6 and 9, which are the nonparametric DGPs. One can clearly see that in DGP 6, it is x_3 that is causing the nonparametricity, since it has a mean $ICR(3)$ of 0.53. Removing this regressor improves the information criterion given the data. The same is true for x_3 and x_{21} in DGP 9, which both

DGP	Variable	$F_N(1.2)$	$ICR_{0.01}$	$ICR_{0.1}$	$E(ICR)$	$ICR_{0.9}$	$ICR_{0.99}$
6	x_3	0.98	0.37	0.38	0.53	0.79	1.40
	x_{11}	0.00	15.79	19.27	25.03	31.33	37.31
9	x_3	0.99	0.75	0.75	0.82	0.94	1.20
	x_{11}	0.00	8.44	9.95	12.56	15.41	18.36
	x_{21}	0.99	0.75	0.75	0.81	0.92	1.18
	x_{29}	0.00	2.15	2.88	4.24	5.73	7.38
	x_{37}	0.00	3.87	5.46	8.82	12.61	17.25

TABLE 7. Distribution of ICRs for DGPs 6 and 9. Notation as in Table 6. Variables that are excluded from the EDGP are in boldface.

have ICRs with a mean of around 0.8. In contrast, removing any of the other regressors has a significant impact on the quality of the model fit. In practice, therefore, one could consider these as the weak regressors.

Therefore, in DGPs 6 and 9, the variables in boldface in Table 7 are excluded from the EDGP. The EDGP are defined as the remaining regressors in each case, see Table 6. For fairness, the results are presented here relative to both the DGP and the EDGP, although it is maintained that the identification of the EDGP is a more reasonable measure of success (a fact reflected by the results of both algorithms, and the original work of HP).